

IN THE CLAIMS

(1) Please rewrite claim 1 as follows:

1 1. (Amended) A method for storing data that has at least some entries with
2 multiple value attributes, comprising the steps of:
#1 3 profiling the data to determine whether the data should be stored in an attribute table
4 or, alternatively, in a merged table and an overflow table; and
5 storing the data optimally based on the profiling step.

(2) Please rewrite claim 5 as follows:

1 5. (Amended) The method as described in Claim 1 wherein a majority of the
#2 2 data is stored in the merged table and a set of additional values for the multiple value
3 attributes are stored in the overflow table.

REMARKS

Claims 1-8 are pending in the Application.

Claims 9-23 have been withdrawn from the Examiner's as being subject to a restriction election requirement.

Claim 1 stands objected to.

Claims 1-8 stand rejected.

I. EXAMINER INTERVIEW SUMMARY

The Applicants and Applicants' attorney appreciate the opportunity to discuss the Application in a telephonic interview with Examiner Newgen on March 4, 2003. In particular, with respect to the rejection under 35 U.S.C. §112, first paragraph, the Examiner suggested adding additional description explaining how the merged table is supposed to be better than the per attribute table. The Applicants express concern that this would be adding

new matter to the Application. Additionally, as discussed further hereinbelow, the Applicants respectfully submit that this does not create a lack of enablement under 35 U.S.C. § 112, first paragraph. The Applicants also pointed to empirical evidence in the Description of the Preferred Embodiment itself with respect to the performance of the merged table and per attribute table. This is also discussed hereinbelow.

II. RESTRICTION UNDER 35 U.S.C. § 121

Restriction to one of the following inventions has been required under 35 § U.S.C. § 121:

Group I, drawn to claims 1-8;

Group II, drawn to claims 9-14;

Group II, drawn to claims 15-18 and 23; and

Group IV, drawn to claims 19-22.

The Applicants provisionally elected Group I, claims 1-8 in a telephonic interview with the Examiner on November 20, 2002. The Applicants hereby affirm the election of Group I, claims 1-8.

III. OBJECTION TO THE CLAIMS

Claim 1 has been objected to because of a typographical error with respect to the clause "data should be in stored in an attribute table." The Applicants, in accordance with the Examiner suggest and have deleted the first occurrence of "in" in the aforementioned clause.

IV. REJECTION UNDER 35 U.S.C. § 112, SECOND PARAGRAPH

Claims 4, 5 and 6 have been rejected under 35 U.S.C. § 112, second paragraph as being indefinite for failing to particularly point out and distinctly claim the subject matter which the Applicants regard as the invention. The Applicants respectfully traverse the rejection of claims 4, 5 and 6.

Claim 4 is directed to the method of claim 1 in which the overflow table is an attribute table. Claim 1 recites a step of profiling data to determine whether the data should be stored in an attribute table, or, alternatively, in a merge table and an overflow table. The Examiner contends that claim 4 does not further limit claim 1, but also alters the scope of the invention in claim 1. The Applicants are unsure as to the rationale for rejecting claim 4 under 35 U.S.C. § 112, second paragraph. Claim 4 includes the express limitation recited therein, in which the overflow table is an attribute table, and incorporates the limitations from claim 1 from which it depends by reference. Therefore, dependent claim necessarily alters the scope of the claim from which it depends. Furthermore, the express limitation from claim 4 further limits the overflow table of claim 1 to be an attribute table. The Applicants respectfully contend that a person of ordinary skill in the art would understand whether his or her overflow table was an attribute table, and therefore, the requirements of 35 U.S.C. § 112, second paragraph are satisfied. *See* MPEP § 2173.02 (stating that the inquiry is to whether the claim as a whole appraises one of ordinary skill in the art of its scope).

Claim 5 is directed to the method of claim 1 in which the majority of the data is stored in the merge table and a small set of additional values for the multiple value attributes are stored in the overflow table. Claim 5 has been rejected under 35 U.S.C. § 112, second paragraph in that small is considered to be vague and indefinite because the term does not indicate a range of data rendering a clear boundary of the claim. Claim 5 has been amended hereinabove to delete the term "small." The amendment to claim 5 is not a narrowing amendment made for a reason relating to the statutory requirements for a patent that will give rise to prosecution estoppel. *See Festo Corp. v. Shoketsu Kinzoku Kogyo Kabushiki Co.*, 122

S. Ct. 1831, 1839-40, 62 U.S.P.Q.2d 1705, 1711-12 (U.S. 2002); 234 F.3d 555, 566, 56 U.S.P.Q.2d 1865, 1870 (Fed. Cir. 2001)).

Claim 6 is directed to the method of claim 1 in which the profiling step parses the data to identifying entries with single value attributes. Claim 6 has been rejected on the ground that the limitation of claim 6 excludes the possibility of entries with multi-value attributes to be stored in an overflow table that was claimed in claim 1. (Paper No. 4, page 5.) The Applicants do not necessarily agree that this precludes a possibility of entries with multi-value attributes to be stored in an overflow table as recited in claim 1. The limitation simply states that the profiling step parses data to identify entries with multiple-value attributes. The Applicants respectfully contend that one of ordinary skill in the art would understand whether his or her method included a profiling step in which the profiling step parses the data to identify entries with single-value attributes. Therefore, the limitation of claim 6 satisfies the essential inquiry under 35 U.S.C. § 112, second paragraph. (See MPEP § 2173.02.)

For at least the aforesaid reasons, the Applicants respectfully contend that claims 4, 5 and 6 satisfy the requirements of 35 U.S.C. § 112, second paragraph. The Applicants respectfully request the Examiner to withdraw the rejection of claims 4, 5 and 6 under 35 U.S.C. § 112, second paragraph.

V. REJECTION UNDER 35 U.S.C. § 112, FIRST PARAGRAPH

Claims 2 and 3 have been rejected under 35 U.S.C. § 112, first paragraph, as containing subject matter which was not described in the Specification in such a way as to enable one skilled in the art to which it pertains, or with which it is most nearly connected, to make and/or use the invention. (Paper No. 4, page 6.) The Applicants respectfully traverse the rejection of claims 2 and 3 under 35 U.S.C. § 112, first paragraph.

Claim 2 is directed to the method of claim 1 in which the entries with single-value attributes are stored in the merged table. The Examiner asserts that "merged table" is merely a naming invention which is used to store entries with single value attributes. As an initial

matter, the Applicants do not understand what is meant by "merely a naming invention." The Examiner further states that the Specification did not further illustrate how one skilled in the art could make use of the invention in terms of how it solves the problem as per attribute table does not perform well for certain operations. (Paper No. 4, page 6.) As noted hereinabove, the Applicants' attorney and the Examiner discussed this issue in a telephonic interview. The Applicants, thus, understand the Examiner to be asserting that claim 2 lacks enablement because the Applicants did not explain how the invention of claim 2 operates to solve the performance problem indicated. However, the Applicants respectfully submit that there is no such requirement under 35 U.S.C. § 112, first paragraph. Indeed, to the contrary, the burden is on the Examiner to provide a reasonable basis to question the enablement provided for in the claimed invention. (MPEP § 2164.04.) It is incumbent upon the Examiner whenever a rejection is made under the enablement to requirement to explain why the truth or accuracy of any statement in the supporting disclosure is doubted and to back up the assertions with acceptable evidence or reasoning which is in consistent with the contested statements. MPEP § 2164.04. Moreover, the foregoing notwithstanding, the Applicants discussed empirical performance results of a prototype. (Detailed Description of the Preferred Embodiment, page 25, line 35 through page 26, line 9.)

Claim 3 is directed to the method of claim 1 in which the entries with multiple value attributes are stored in the overflow table. Claim 3 has been rejected on the same rationale in which the "overflow table" is another name describing multiple value attributes. (Paper No. 4, page 6.) Again, the Applicants are unsure by what is meant by the overflow table being another name describing multiple value attributes. As discussed in conjunction with claim 2, the Applicants respectfully assert that claim 3 has not been shown to lack description in the Specification in such a way as to enable one skilled in the relevant art to make or use the invention thereof.

VI. REJECTION UNDER 35 U.S.C. § 103

Claims 1-4 and 6-8 have been rejected under 35 U.S.C. § 103 as being unpatentable over *Morgenstern*, U.S. Patent No. 5,970,490 in view of *Gioielli, et al.*, U.S. Patent

No. 5,485,610 ("*Gioielli*"). The Applicants respectfully traverse the rejection of claims 1-4 and 6-8 under 35 U.S.C. § 103.

Claim 1 is directed to a method for storing data that has at least some entries with multiple value attributes. The method includes the steps of profiling the data to determine whether the data should be stored in an attribute table, or, alternatively, in a merged table and an overflow table, and storing the data optimally based on the profiling step. *Morgenstern* allegedly teaches the limitations of claim 1 but for the storing of the data optimally based on a profiling step. The Applicants respectfully disagree.

Morgenstern is directed to integration platforms for heterogeneous databases, and in particular to a method for heterogeneous data using an interoperability assistant module with specifications for transforming the data into a common intermediate representation of the data using the specifications and creating an information bridge with the interoperability assistant module through a process of program generation. (*Morgenstern*, column 1, lines 10-17.) Thus, *Morgenstern* is not related to methods for storing data in a directory service backing store.

With respect to the limitations of claim 1, *Morgenstern* is asserted to teach the limitation with regard to profiling the data to determine whether the data should be stored in a attribute table, or, alternatively, in a merged table and an overflow table. In particular, this limitation is purportedly taught in disclosing that if the input high level data structure specifications refer to structured files, such as design files, for specially formatted data, then the recognizer generator also produces yacc-like code that the parsing tool can process. (Paper No. 4, page 7) (citing *Morgenstern*, column 8, lines 38-47).¹ The teaching relied upon further discloses that the parsing tool generates the actual recognition/parsing module that is incorporated into the information bridge. (*Id.*) The teaching in *Morgenstern* is further relied upon in disclosing the aforementioned step of claim 1 discloses that the sequence of values for different leaf nodes is defined as a logical 'tuple' since there is a single value in each position at any one time and that a tuple is defined as the sequence of data values, one

per leaf node, that fall under the tree scope of a selected node which is referred to as an 'anchor' node. (Paper No. 4, page 7) (citing *Morgenstern*, column 32, lines 43-45 for the proposition that *Morgenstern* teaches determining whether data should be stored in a merged table). The Examiner also relies on teaching in *Morgenstern* disclosing that the Semantic Metadata Description Language is a frame-based representation where a MetaFrame contains potentially multiple attribute value specifications wherein each attribute may have subattributes, and these two may have subattributes if needed thereby creating a potentially hierarchical array of descriptors such as 'units' and 'precision' information for length, weight or other measures. (Paper No. 4, page 7) (citing *Morgenstern*, column 38, lines 50-55 for the proposition that *Morgenstern* teaches an overflow table). Plainly, by the express terms of the teaching in *Morgenstern*, *Morgenstern* has not been shown to teach or suggest a step of profiling the data should be stored in an attribute table or, alternatively, in a merged table and an overflow table.

With respect to the step of storing the data optimally based on the profiling step, the Examiner relies on *Gioelli* to supply the limitations admittedly missing in *Morgenstern*. (Paper no. 4, page 7.) *Gioelli* is directed to a physical database designer embodied in computer software that generates a physical database design. (*Gioelli*, column 1, lines 65-67.) The physical structure of a database is the particular storage characteristics of a collection of interrelated data stored as one or more stored records (*Gioelli*, column 1, lines 39-41.) A stored record is the basic unit of storage which corresponds to one logical record and contains all pointers, record lengths and other identifiers necessary to represent pieces of data. (*Gioelli*, column 1, lines 42-45.) Aspects of the physical structure of a database include the location and size of files, the access methods and keys to database records and the placement of records within the database. (*Gioelli*, column 1, lines 45-52.) Thus, in sum, *Gioelli* is directed to mechanisms for automating the design of a physical structure of a database, that is, the particular storage characteristics of data stored as one or more stored records. With respect to the disclosure in *Gioelli* purported to teach the

¹ Yacc is an acronym referring to "Yet Another Compiler-Compiler." Yacc is a tool for imposing structure on the input to a computer program. *Id.* There is nothing in disclosing yacc-like code that teaches the limitation of claim 4. STEPHEN JOHNSON, YACC: YET ANOTHER COMPILER-COMPILER, <http://dinosaur.compilertools.net/yacc> at 1.

limitation of storing the data optimally based on that profiling step, the Examiner relies on teaching in *Gioelli* discussing run time parameters (used to unload data from an old database and reload it in a new database) and creation parameters (used to implement the physical database structure) contained in a command procedure, that is, a series of Digital Command Language commands that the user can use to implement the improved database; this teaching also discloses that when the command procedure is run, the data entities from an existing database are unloaded, the physical design of the database is optimized, and the data entities reloaded into the new database. (Paper No. 4, page 7) (citing *Gioelli*, column 15, lines 13-28). Thus, the teaching relied upon as disclosing the step of optimally storing data based on the profiling step refers to optimization of a physical design of a database, not the optimal storing data based on a step of profiling the data.

Thus, for at least these reasons, the Applicants respectfully contend that neither *Morgenstern* nor *Gioelli*, alone or in combination teach or suggest all of the limitations of claim 1.

With respect to a motivation for combining the references, it is asserted that it would have been obvious to incorporate the teaching in *Gioelli* into *Morgenstern* to optimize the physical design of the database. (Paper No. 4, page 7.) However, *Morgenstern*, as previously discussed, is directed to mechanisms for processing heterogeneous data, and in particular, mechanisms for transforming data into a common intermediate representation of the data. (See e.g., *Morgenstern*, column 1, lines 10-17.) It is by no means transparent that the teaching of *Gioelli* with respect to optimization may be incorporated into *Morgenstern*; the likelihood of success must be demonstrated by objective evidence found in the references themselves. MPEP § 2143. That notwithstanding, the combination of *Morgenstern* and *Gioelli*, as discussed hereinabove, do not teach or suggest the invention of claim 1.

Consequently, for at least these reasons, the Applicants respectfully assert that a *prima facie* showing of obviousness has not been made with respect to claim 1. Therefore, claim 1 is allowable under 35 U.S.C. § 103 over *Morgenstern* and *Gioelli*, and the Applicants respectfully request the rejection of claim 1 under 35 U.S.C. § 103 be withdrawn.

Claim 2 depends from claim 1 and recites the method thereof in which the entries with single value attributes are stored in the merged table. *Morgenstern* is purported to teach the limitation of claim 2. (Paper No. 4, page 7.) The teaching in *Morgenstern* discusses a relational application programming interface (API) which present data in terms of named arrays (relation tables) that contain flat tuples, with one data value per column and each tuple, and with a number of columns fixed per named array. (See *Morgenstern*, column 1, lines 47-52.) In other words, the teaching in *Morgenstern* is directed to an array of data in which the rows (flat tuples) contain one data value per column and the number of columns is fixed for each named array. There is nothing in this teaching of *Morgenstern* that discusses storing single value attributes in a merged table. Furthermore, as the Applicants have discussed, both single value attributes and multi-value attributes may be stored in relation tables. (Description of the Related Art, page 3, lines 14 through page 4, line 2.) Thus, the Applicants respectfully contend that *Morgenstern* has not been shown to teach or suggest the limitations of claim 2. Consequently, for this reason and those discussed in conjunction with claim 1 above, the Applicants respectfully assert that a *prima facie* showing of obviousness has not been made with respect to claim 2 and claim 2 is allowable under 35 U.S.C. § 103 over *Morgenstern* and *Gioelli*.

Claim 3 is directed to the method of claim 1 in which entries with multiple-value attributes are stored in the overflow table. *Morgenstern* is alleged to teach the limitation of claim 3 in disclosing the representation of Metadata as the Semantic Metadata Description Language, a frame-based representation where a MetaFrame contains potentially multiple attribute value specifications. (Paper No. 4, page 7) (citing *Morgenstern*, column 38, lines 47-57). Note that multiple attribute value specifications are not the same as multi-valued attributes. The teaching additionally discusses that each attribute may have subattributes and these too may have subattributes, creating a potentially hierarchical array of descriptors. (*Morgenstern*, column 38, lines 52-57.) The Examiner also points to a single sentence in *Morgenstern* stating that multiple values for an attribute are allowed, where ordering is important. (Paper No. 4, page 7) (citing *Morgenstern*, column 40, lines 27-28). Note however that this teaching makes no reference to the storage of multiple-value attributes in an overflow table. Thus, the Applicants respectfully contend that *Morgenstern*

has not been shown to teach or suggest all of the limitations of claim 3, and thus neither *Morgenstern*, alone or in combination with *Gioelli* teach or suggest all of the limitations of claim 3. Therefore, claim 3 is allowable under 35 U.S.C. § 103 over *Morgenstern* and *Gioelli*.

Claim 4 is directed to the method of claim 1 in which the overflow table is an attribute table. *Morgenstern* is alleged to teach the limitation of claim 4 in disclosing the metadata representation discussed above in conjunction with claim 3. (Paper No. 7, page 8) (citing *Morgenstern*, column 38, lines 47-57). Again, by the plain teaching of the referred to disclosure in *Morgenstern*, *Morgenstern* has not been shown to teach or suggest all of the limitation of claim 4. Therefore, for at least this reason and those discussed hereinabove in conjunction with claim 1, the Applicants respectfully contend that a *prima facie* showing of obviousness has not been made with respect to claim 4. Therefore, claim 4 is allowable under 35 U.S.C. § 103 over *Morgenstern* and *Gioelli*.

Claim 6 is directed to the method of claim 1 in which the profiling step parses the data to identify entries with single value attributes. Claim 6 has been rejected over teaching in *Morgenstern* discussed hereinabove in conjunction with claim 1, namely teaching directed to the high level data structure specifications referring to structured files such as design files, for specially formatted data, and teaching directed to a sequence of values for different leaf nodes being defined as a logical tuple. (Paper No. 4, page 8) (citing *Morgenstern*, column 8, lines 39-42 and column 32, lines 43-45). The Examiner also refers to teaching in *Morgenstern* discussing a merge operator combining different input tokens from the left side into one sequence which is then passed to the body of the transformation rule, the result being a single instance which is given the label on the right side of the rule. (Paper No. 4, page 8) (citing *Morgenstern*, column 19, lines 40-45). The Applicants are unsure as to the relevance of this teaching to the limitations in claim 6. Nevertheless, by the plain teaching of the disclosure referred to in conjunction with claim 6, *Morgenstern* has not been demonstrated to teach or suggest the limitations of claim 6. For this reason, and those discussed hereinabove in conjunction with claim 1, the Applicants respectfully contend that

a *prima facie* showing of obviousness has not been made with respect to claim 6. Therefore, claim 6 is allowable under 35 U.S.C. § 103 over *Morgenstern* and *Gioelli*.

Claim 7 is directed to the method of claim 1 in which the profiling step parses the data to identify given operations that are performed on the data once stored. Claim 7 has been rejected on teaching in *Morgenstern* discussing parsers as data access functions associated with each uniform region. (*Morgenstern*, column 26, line 53 through column 27, line 26.) This teaching discusses, for example, that different uniform regions will have different instances of the same parser or object type when these regions have the same type of data source and same type of data representations, for example, objects in an object database. (*Morgenstern*, column 27, lines 1-4.) *Morgenstern* further teaches that a given uniform region is characterized by a data source, a parser type and a parser object instance. (*Morgenstern*, column 27, lines 9-12.) *Morgenstern* additionally teaches that the process of determining uniform regions in associated parsers for the logical structure diagram tree serves to cover all terminal nodes with distinct parsers. (*Morgenstern*, column 27, lines 13-18.) *Morgenstern* also discloses that all nodes which are outside some uniform region are designated as parser controller nodes and these are parents of either uniform regions and/or other parser controller nodes. (*Morgenstern*, column 27, lines 19-26.) Thus, the teaching referred to, by the plain terms thereof, does not disclose a profiling step in which the profiling step parses the data to identify given operations that are performed on the data once stored. Because neither *Morgenstern* nor *Morgenstern* in combination with *Gioelli* teach or suggest all of the limitations of claim 7, nor is there a motivation for modifying or combining *Morgenstern* and *Gioelli* to make the claimed invention, the Applicants respectfully contend that a *prima facie* showing of obviousness has not been made with respect to claim 7. Therefore, claim 7 is allowable under 35 U.S.C. § 103 over *Morgenstern* and *Gioelli*.

Claim 8 is directed to the method of claim 1 in which the data is stored in a relational database backing store. Claim 8 has been rejected on teaching in *Morgenstern* that teaches that a given schema tree can map back to data that is stored in very different ways, like a flat file, a web page, a relational database or an object oriented database, but share the same logical structure, by the use of different annotations for the input and output. (See *Morgenstern*, column 24, lines 45-49.) The Applicants respectfully contend that there is

nothing in this teaching in *Morgenstern* that suggests storing the data as recited in claim 1 in a relational database backing store. Nevertheless, claim 8 incorporates all of the limitations of claim 1 from which it depends, and as discussed hereinabove, neither *Morgenstern* nor *Morgenstern* and *Gioelli* alone or in combination teach or suggest all of the limitations of claim 8. Therefore, the Applicants respectfully contend that a *prima facie* showing of obviousness has not been made with respect to claim 8, and claim 8 is allowable under 35 U.S.C. § 103 over *Morgenstern* and *Gioelli*.

VII. CONCLUSION


As a result of the foregoing, it is asserted by Applicants that the remaining claims in the Application are in condition for allowance, and respectfully request an early allowance of such claims.

Applicants respectfully request that the Examiner call Applicants' attorney at the below listed number if the Examiner believes that such a discussion would be helpful in resolving any remaining problems.

Respectfully submitted,

WINSTEAD SECHREST & MINICK P.C.

Attorney for Applicants

By: 
Barry S. Newberger
Reg. No. 41,527

5400 Renaissance Tower
1201 Elm Street
Dallas, Texas 75270-2199
(512) 370-2808

VERSION TO SHOW CHANGES MADE**IN THE CLAIMS**

(1) Claim 1 has been rewritten as follows:

- 1 1 . (Amended) A method for storing data that has at least some entries with multiple
2 value attributes, comprising the steps of:
3 profiling the data to determine whether the data should be [in] stored in an attribute
4 table or, alternatively, in a merged table and an overflow table; and
5 storing the data optimally based on the profiling step.

(2) Claim 5 has been rewritten as follows:

- 1 5. (Amended) The method as described in Claim 1 wherein a majority of the data
2 is stored in the merged table and a [small] set of additional values for the multiple value
3 attributes are stored in the overflow table.

AUSTIN_1\209724\1
7047-P398US